

A Computationally Efficient Implementation of Fictitious Play for Large-Scale Games

BRIAN SWENSON^{†*}, SOUMMYA KAR[†] AND JOÃO XAVIER^{*}

Abstract

The paper is concerned with distributed learning and optimization in large-scale settings. The well-known Fictitious Play (FP) algorithm has been shown to achieve Nash equilibrium learning in certain classes of multi-agent games. However, FP can be computationally difficult to implement when the number of players is large. Sampled FP is a variant of FP that mitigates the computational difficulties arising in FP by using a Monte-Carlo (i.e., sampling-based) approach. The Sampled FP algorithm has been studied both as a tool for distributed learning and as an optimization heuristic for large-scale problems. Despite its computational advantages, a shortcoming of Sampled FP is that the number of samples that must be drawn in each round of the algorithm grows without bound (on the order of \sqrt{t} , where t is the round of the repeated play). In this paper we propose Computationally Efficient Sampled FP (CESFP)—a variant of Sampled FP in which only one sample need be drawn each round of the algorithm (a substantial reduction from $O(\sqrt{t})$ samples per round, as required in Sampled FP). CESFP operates using a stochastic-approximation type rule to estimate the expected utility from round to round. It is proven that the CESFP algorithm achieves Nash equilibrium learning in the same sense as classical FP and Sampled FP. Simulation results suggest that the convergence rate of CESFP (in terms of repeated-play iterations) is similar to that of Sampled FP.

I. INTRODUCTION

A game-theoretic learning algorithm is an adaptive multi-agent procedure which can enable a system of interacting agents to achieve desirable global behavior using local (agent-based)

The work was partially supported by the FCT project FCT [UID/EEA/50009/2013] through the Carnegie-Mellon/Portugal Program managed by ICTI from FCT and by FCT Grant CMU-PT/SIA/0026/2009, and was partially supported by NSF grant ECCS-1306128.

[†]Department of Electrical and Computer Engineering, Carnegie Mellon University, Pittsburgh, PA 15213, USA (brian-swe@andrew.cmu.edu and soumyak@andrew.cmu.edu).

^{*}Institute for Systems and Robotics (ISR/IST), LARSyS, Instituto Superior Técnico, Portugal (jxavier@isr.ist.utl.pt).

control laws. Such algorithms have a wide range of applications in distributed control [1]–[5] and large-scale optimization [6]–[10].

Fictitious Play (FP) [11] is an archetypal game-theoretic learning algorithm that has received much attention over the years due to its intuitively simple nature and proven convergence results for certain important classes of games (see, for example, [12], [13] and references therein). Of particular interest are results demonstrating that FP leads players¹ to learn equilibrium strategies in potential games—a class of multi-agent games in which there may be an arbitrarily large number of players [14], [15].

However, such convergence results tend to be of limited practical value due to computational difficulties that may arise when implementing FP in large games. In particular, in each stage of the FP algorithm, each player i must compute the expected (mixed) utility for each of her actions given her current beliefs regarding opponents’ strategies. Evaluating this expected utility—the domain of which is an $(n - 1)$ -dimensional probability simplex—is a problem whose complexity in general scales exponentially in terms of the number of players, n .

The main focus of the present paper is the presentation of a variant of FP that might be more practical to implement in certain large-scale settings. In particular, we consider a practical method for mitigating computational complexity using a Monte-Carlo type approach.

Sampled FP [9], [10], [16]–[19] introduced the idea of mitigating complexity in FP using a Monte-Carlo (i.e., sampling-based) approach. At each iteration of the Sampled FP algorithm, players approximate the expected utility by drawing several samples from an underlying probability distribution. Players then myopically choose an “optimal” next-stage action using the approximated utility as a surrogate for the true expected utility. The work [9] showed that, as long as the number of samples drawn each round grows sufficiently quickly, players learn an equilibrium in the same sense as FP, almost surely (a.s.).

In essence, Sampled FP achieves a mitigation in complexity by avoiding any direct evaluation of the expected utility. However, Sampled FP has a notable shortcoming: In order to guarantee learning is achieved, the number of samples that must be drawn in each iteration (i.e., round) of the algorithm grows without bound (on the order of \sqrt{t} samples per round, where t is the current round of the repeated play algorithm).

¹We use the terms agent and player interchangeably throughout the paper.

In this paper, we propose a variant of Sampled FP—which we call Computationally Efficient Sampled FP (CESFP)—in which only one sample need be drawn each round of the repeated play process. CESFP achieves the same fundamental computational advantage as Sampled FP (i.e., direct evaluation of the expected utility is avoided) but does so by drawing only one sample per round (rather than the $O(\sqrt{t})$ samples per round, required in Sampled FP).

Intuitively, the reduction in the required number of per-round samples is accomplished by treating the expected utility process as *quasi static*. Such treatment is possible due to the diminishing incremental step size in the the expected utility process. In CESFP, the sample data gathered in the current round of repeated play is recursively combined with sample data from previous rounds using a stochastic-approximation-type estimation rule. This may be contrasted with Sampled FP where, in each round of the repeated play, data gathered from sampling in the previous round is wholly discarded and a fresh set of samples is gathered to approximate the expected utility for the upcoming round. (See Section V-C for more details.)

Due to the improved efficiency in information handling, CESFP is able to achieve convergence at a rate similar to that of sampled FP (in terms of repeated-play iterations) despite drawing far fewer samples per-iteration. (See Section VI for more details.)

The main contribution of the paper is the presentation of the CESFP algorithm and proof of convergence of the algorithm in terms of empirical frequency to the set of Nash equilibria (a.s.). The proof relies on showing that the CESFP process may be seen as a Generalized Weakened FP process as studied in [20].

CESFP may be applicable as a computationally efficient variant of FP in a variety of settings including large-scale optimization [9], [18], dynamic programming [16], [17], traffic routing [10], and cognitive radio [4], [21], [22], and learning in Markov decision processes [19]. CESFP may also be used as a general tool for distributed learning [23] or control [1].

Related works have studied approaches for mitigating computational issues arising in large-scale implementations of FP. Joint Strategy FP (JSFP) [24] studies a variant of FP in which players update a utility estimate using a computationally-simple recursive procedure and choose next-stage actions using a best-response rule combined with an inertial term. JSFP is shown to converge to pure strategy Nash equilibria (NE) in ordinal potential games but is fundamentally different from CESFP (and FP) in that the tracked utility corresponds to an empirical distribution

taken over joint actions,² and convergence may only occur at pure NE.

Payoff-based learning algorithms—including those based on FP [20], [25]–[27] and otherwise [28]–[30]—tend to be computationally simple in large games and have the further advantage that they do not require players to have any knowledge of the game’s utility structure. However, such algorithms implicitly assume players have access to instantaneous payoff information and may not be applicable in settings where this information is costly to obtain, delayed, or otherwise unavailable.

For example, in a follow-up paper [23] we consider an application in which CESFP is implemented in a network-based setting in which all inter-agent communication is restricted to a preassigned (possibly sparse) communication graph [31]. Instantaneous payoff information can be difficult to obtain in such a setting, particularly in the case that the utility corresponds to a non-local welfare-type utility function, which may not be physically measurable at any single agent. Furthermore, there are circumstances in which each round of physical game play can incur an exogenous cost. In such cases it may be preferable to supplement payoff-based learning—which depends on interaction in the physical environment—with forms of model-based learning.

The remainder of the paper is organized as follows. Section II sets up the notation to be used in the subsequent development. Section III reviews classical FP. Section IV reviews Sampled FP. Section V presents the CESFP algorithm, states the main convergence result for CESFP, and proves the result. Section VI presents a simulation example comparing Sampled FP and CESFP. Section VII provides concluding remarks.

II. PRELIMINARIES

A game in normal form is represented by the tuple $\Gamma := (N, (Y_i, u_i)_{i \in N})$, where $N = \{1, \dots, n\}$ denotes the set of players, Y_i denotes the finite set of actions available to player i , and $u_i : \prod_{i \in N} Y_i \rightarrow \mathbb{R}$ denotes the utility function of player i . Denote by $Y := \prod_{i \in N} Y_i$ the joint action space.

In order to guarantee the existence of Nash equilibria it is necessary to consider the mixed-extension of Γ in which players are permitted to play probabilistic strategies. Let $m_i := |Y_i|$ be

²In FP, Sampled FP, and CESFP, players best respond to the product of marginal empirical distributions (or an estimate thereof) which implicitly presumes a form of independence among opponents’ strategies. Tracking and responding to the empirical distribution of joint actions, as in JSFP, fundamentally alters the dynamics of classical FP. CESFP achieves computationally efficiency while preserving the basic dynamical structure of FP.

the cardinality of the action space of player i , and let $\Delta_i := \{p \in \mathbb{R}^{m_i} : \sum_{k=1}^{m_i} p(k) = 1, p(k) \geq 0, \forall k\}$ denote the set of mixed strategies available to player i —note that a mixed strategy is a probability distribution over the action space of player i . Denote by $\Delta^n := \prod_{i \in N} \Delta_i$ the set of joint mixed strategies. When convenient, we represent a mixed strategy $p \in \Delta^n$ by $p = (p_i, p_{-i})$, where p_i denotes the marginal strategy of player i and p_{-i} is a $(n-1)$ -tuple containing the marginal strategies of the other players.

In the context of mixed strategies, we often wish to retain the notion of playing a single deterministic action. For this purpose, let $A_i := \{e_1, \dots, e_{m_i}\}$ denote the set of “pure strategies” of player i , where e_j is the j -th canonical vector containing a 1 at position j and zeros otherwise. Note that there is a one-to-one correspondence between a player’s action set Y_i and the player’s set of pure strategies $A_i \subset \Delta_i$.

The mixed utility function of player i is given by

$$U_i(p) := \sum_{y \in Y} u_i(y) p_1(y) \dots p_n(y) \quad (1)$$

where $U_i : \Delta^n \rightarrow \mathbb{R}$. Note that the mixed utility $U_i(p)$ may be interpreted as the expected utility of $u_i(y)$ given that players’ (marginal) mixed strategies p_i are independent.

The set of Nash equilibria is given by $NE := \{p \in \Delta^n : U_i(p_i, p_{-i}) \geq U_i(p'_i, p_{-i}), \forall p'_i \in \Delta_i, \forall i \in N\}$. The distance of a distribution $p \in \Delta^n$ from a set $S \subset \Delta^n$ is given by $d(p, S) = \inf\{\|p - p'\| : p' \in S\}$. Throughout the paper $\|\cdot\|$ denotes the standard \mathcal{L}_2 Euclidean norm unless otherwise specified.

Throughout, we assume there exists a probability space $(\Omega, \mathcal{F}, \mathbb{P})$ rich enough to carry out the construction of the various random variables required in the paper. For a random object X defined on a measurable space (Ω, \mathcal{F}) , let $\sigma(X)$ denote the σ -algebra generated by X [32]. As a matter of convention, all equalities and inequalities involving random objects are to be interpreted almost surely (a.s.) with respect to the underlying probability measure, unless otherwise stated.

A. Repeated Play

The learning algorithms considered in this paper all assume the following format of repeated play. Let a normal form game Γ be fixed. Let players repeatedly face off in the game Γ , and for $t \in \{1, 2, \dots\}$, let $a_i(t) \in A_i$ denote the action played by player i in round t . Let the n -tuple $a(t) = (a_1(t), \dots, a_n(t))$ denote the joint action at time t .

Let the empirical history distribution (or empirical distribution) of player i be given by³ $q_i(t) := \frac{1}{t} \sum_{s=1}^t a_i(s)$, and let the joint empirical distribution be given by the n -tuple $q(t) = (q_1(t), \dots, q_n(t))$.

III. FICTITIOUS PLAY

A FP process may be intuitively described as follows: A finite set of n agents engage in repeated play of some fixed normal form game. Each round of the repeated play, each agent i plays an action that is myopically optimal under the (naive) assumption that all opponents are playing according to time-invariant and statistically independent strategies. In particular, under this assumption, each player i believes that the empirical distribution $q_{-i}(t)$ of opponents' play is an accurate representation of opponents' (supposedly time-invariant) strategies and chooses a next-stage action that optimizes their utility given this belief.

A formal description of the FP algorithm is given below.

A. FP algorithm

Initialize

(i) Each player i chooses an arbitrary initial action $a_i(1) \in A_i$. The empirical distribution is initialized as $q_i(1) = a_i(1)$, $\forall i$.

Iterate ($t \geq 1$)

(ii) Each player i chooses her next-stage action as a best response to the current empirical distribution of opponents' play:

$$a_i(t+1) \in \arg \max_{\alpha_i \in A_i} U_i(\alpha_i, q_{-i}(t)). \quad (2)$$

(iii) For each player i , the empirical distribution is updated to reflect the action just taken, $q_i(t+1) = \frac{1}{t+1} \sum_{s=1}^{t+1} a_i(s)$, or equivalently in recursive form:

$$q_i(t+1) = q_i(t) + \frac{1}{t+1} (a_i(t+1) - q_i(t)).$$

³Note that each $a_i(t) \in A_i$ is a delta distribution, and thus the empirical distribution $q_i(t)$ is a normalized histogram of the action choices of player i .

B. Discussion

Various works have established results characterizing classes of games in which FP does (or does not) lead players to learn NE strategies (see monographs [12], [13]). Of particular relevance to the large-scale setting is a class of multi-agent games known as potential games [15]. A game $\Gamma = (N, (Y_i, u_i(\cdot))_{i \in N})$ is said to be a potential game if there exists a potential function $\phi : Y^n \rightarrow \mathbb{R}$ such that for all $i \in N$, and all $y_{-i} \in Y_{-i}$

$$u_i(y_i, y_{-i}) - u_i(x_i, y_{-i}) = \phi(y_i, y_{-i}) - \phi(x_i, y_{-i}), \quad \forall y_i, x_i \in Y_i.$$

Intuitively, the existence of a potential function means that all player's utility functions are aligned in such way that players share a common underlying objective. It has been shown [14], [15] that if Γ is a potential game, then FP leads players to learn NE strategies in the sense that $\lim_{t \rightarrow \infty} d(q(t), NE) = 0$.

C. Computational Complexity in large-scale FP

While FP is theoretically proven to achieve NE learning in potential games, it can be computationally difficult to implement when the number of players is large. In particular, note that in order to choose a next-stage action (see (2)) player i must compute the mixed utility $U_i(\alpha_i, q_{-i}(t))$, $\forall \alpha_i \in A_i$. Recalling the definition of mixed utility (1), this is equivalent to computing an expected value over an $(n - 1)$ -dimensional probability simplex. In general, the complexity of this computation grows exponentially in terms of the number of players.

IV. SAMPLED FP

In order to mitigate the problem of computational complexity in FP, [9] proposed Sampled FP. In Sampled FP, players use a Monte-Carlo approach to avoid direct evaluation of the mixed utility when choosing a next-stage action.

In particular, for each $\alpha_i \in A_i$, let $\hat{U}_i(\alpha_i, t)$ denote an estimate that player i forms of the mixed utility $U_i(\alpha_i, q_{-i}(t))$. Each round of play, for each player i , several “test actions” are drawn as random samples from opponents’ joint empirical distribution $q_{-i}(t)$. For each action $\alpha_i \in A_i$, player i computes the average utility the action α_i would generate given the randomly sampled “test actions.” Player i then chooses a next-stage action that is myopically optimal using

the estimated utility as a surrogate for the true mixed utility in (2). So long as the number of samples increases sufficiently quickly, it can be shown that Sampled FP leads players to learn NE strategies.

Formally, let k_t denote the number of samples drawn in round t . The following assumption on k_t is sufficient to ensure learning is achieved.

A. 1. *The number of samples drawn in round t satisfies $k_t = \lceil Ct^\gamma \rceil$ where $\gamma > 1/2$ and $C > 0$.*

The Sampled FP algorithm is outlined below.

A. Sampled FP Algorithm

Initialize

(i) Each player i chooses an arbitrary initial action $a_i(1) \in A_i$. The empirical distribution is initialized as $q_i(1) = a_i(1)$, $\forall i$.

Iterate ($t \geq 1$)

(ii) k_t “test actions” are drawn as random samples from the joint empirical distribution $q(t)$; let $\tilde{a}^s(t)$ denote the s -th random sample drawn in round t . Player i estimates the utility of each of her actions $\alpha_i \in A_i$ by⁴

$$\hat{U}_i(\alpha_i, t) = \frac{1}{k_t} \sum_{s=1}^{k_t} U_i(\alpha_i, \tilde{a}_{-i}^s(t)). \quad (3)$$

(iii) Each player i chooses a next-stage action that is a best response given her estimate of the mixed utility:

$$a_i(t+1) \in \arg \max_{\alpha_i \in A_i} \hat{U}_i(\alpha_i, t). \quad (4)$$

(iv) The empirical distribution for each player i is updated recursively to account for the action just taken:

$$q_i(t+1) = q_i(t) + \frac{1}{t+1} (a_i(t+1) - q_i(t)).$$

B. Convergence Result

The following result ([9], Theorem 5) shows that under **A.1**, Sampled FP achieves NE learning in the same sense as classical FP for a special class of potential games known as

⁴Since $\tilde{a}_{-i}(t)$ is a pure strategy, the evaluation of the utility is relatively simple. Note also that it is not necessary to draw a separate sequence of “test actions” $\{\tilde{a}^s(t)\}_{s=1}^{k_t}$ for each player i , although, doing so does not affect the convergence result.

identical interests games [14], in which all players use an identical utility function; i.e., $u_i(y) = u_j(y)$, $\forall y \in Y$, $\forall i, j \in N$.

Theorem 1 (Theorem 5 in [9]). *Let Γ be a finite game in strategic form with identical payoffs. Then, any Sampled FP process with sample sizes satisfying **A.1** converges in beliefs to equilibrium with probability 1. That is, $d(q(t), NE) \rightarrow 0$ almost surely as $t \rightarrow \infty$.*

V. COMPUTATIONALLY EFFICIENT SAMPLED FP

While Sampled FP does obtain computational savings when compared to classical FP, it may be considered unsatisfactory in the sense that it requires players to draw a number of samples each round that grows without bound (see **A.1**). In this section, we present an adaptation of Sampled FP in which only one sample need be drawn each round of the repeated play.

A. Algorithm Setup

In CESFP, players form an estimate of the mixed utility using a recursive stochastic-approximation-type rule. Similar to Sampled FP, let $\hat{U}_i(\alpha_i, t)$ be the estimate which player i maintains of the mixed utility $U_i(\alpha_i, q_{-i}(t))$ for each of her actions $\alpha_i \in A_i$. Let $\{\rho(t)\}_{t \geq 1}$ be a deterministic sequence of weights to be used in the stochastic-approximation-type procedure, and assume:

A. 2. *The sequence $\{\rho(t)\}_{t \geq 1}$ is such that $0 < \rho(t) \leq 1$, $\sum_{t \geq 1} (\rho(t))^2 < \infty$, and $\lim_{t \rightarrow \infty} \frac{1}{t\rho(t)} = 0$.*

Note that, by Lemma 3, **A.2** implies that $\sum_{t \geq 1} \rho(t) = \infty$. The Computationally Efficient Sampled FP algorithm is outlined below.

B. Computationally Efficient Sampled FP Algorithm

Initialize

(i) For each $i \in N$, let $a_i(1) \in A_i$ be arbitrary. Initialize the empirical distribution as $q_i(1) = a_i(1)$, $\forall i$, and initialize the utility estimate as $\hat{U}(\alpha_i, 0) = 0$, $\forall \alpha_i \in A_i$, $\forall i$.

Iterate ($t \geq 1$)

(ii) A single “test action” $a^*(t)$ is drawn as a (statistically independent) random sample from

the distribution $q(t)$, and each player i updates the estimate $\widehat{U}_i(\alpha_i, t)$ for each action $\alpha_i \in A_i$ according to the recursion,⁵

$$\widehat{U}_i(\alpha_i, t) = (1 - \rho(t))\widehat{U}_i(\alpha_i, t-1) + \rho(t)U_i(\alpha_i, a_{-i}^*(t)). \quad (5)$$

(iii) Each player i chooses a next-stage action using the rule (cf. (2), (4)):

$$a_i(t+1) \in \arg \max_{\alpha_i \in A_i} \widehat{U}_i(\alpha_i, t). \quad (6)$$

(iv) The empirical distribution for each player i is updated to reflect the action just taken:

$$q_i(t+1) = q_i(t) + \frac{1}{t+1}(a_i(t+1) - q_i(t)). \quad (7)$$

C. Discussion

The main difference between Sampled FP and CESFP is the manner in which players form estimates of the mixed utility sequence $\{U(\alpha_i, q_{-i}(t))\}_{t \geq 1}$, $\forall \alpha_i \in A_i$. In Sampled FP, players' estimates (see (3)) “start afresh” each round of the repeated play—information gathered from sampling in the previous round is discarded, and players draw roughly \sqrt{t} new samples in order to form an estimate of the utility for the current round.

This may be considered an inefficient use of information, since the mixed utility only changes slightly from one round to the next. In particular, note that the mixed utility $U_i(\alpha_i, \cdot)$ is Lipschitz continuous with some Lipschitz constant K , and the increment of the empirical distribution (7) is bounded as $\|q(t) - q(t-1)\| \leq \frac{M}{t}$ for some constant $M > 0$. Thus, the increment in the mixed utility is bounded as

$$|U_i(\alpha_i, q_{-i}(t)) - U_i(\alpha_i, q_{-i}(t-1))| \leq KM/t. \quad (8)$$

Intuitively speaking, this means that if one has an accurate estimate of the mixed utility $U_i(\alpha_i, q_{-i}(t-1))$ in round $(t-1)$, then it is wasteful to wholly discard this information when forming an estimate of $U_i(\alpha_i, q_{-i}(t))$. The CESFP estimation rule leverages the diminishing increment property (8) in order to form an accurate estimate using only one sample per round.

⁵Since $a_{-i}^*(t)$ is a pure strategy, the evaluation of the utility is relatively simple. Also note that it is not necessary for each player i to draw a separate “test action” $a_{-i}^*(t+1)$; although, if desired (for example, in a distributed setting [23]), doing so does not affect the convergence result.

Effectively, the Sampled FP estimation rule treats $\{U_i(\alpha_i, q_{-i}(t))\}_{t \geq 1}$ as if it were arbitrarily generated from one round to the next—drawing a completely new set of t^γ , $\gamma > 1/2$ samples to estimate each $U_i(\alpha_i, q_{-i}(t))$. The CESFP estimation rule, on the other hand, treats $\{U_i(\alpha_i, q_{-i}(t))\}_{t \geq 1}$ as if it were *quasi static*, drawing one sample per round, and taking a type of average over time.⁶ Because of this, despite drawing only one sample per round, the CESFP estimate of $U_i(\alpha_i, q_{-i}(t))$ *effectively* utilizes information from t samples, while the Sampled FP estimate utilizes information from (only) t^γ , $\gamma > 1/2$ samples. In practice, the CESFP and Sampled FP estimation rules tend to reduce estimation error at comparable rates. See Section VI for more details.

D. Main Result

The following theorem states that CESFP achieves learning in the same sense as classical FP (and Sampled FP). The result is stated for a slightly broader classes of games than previously discussed, including two-player zero-sum games [12], and generic $2 \times m$ games [35].

Theorem 2. *Let Γ be a potential game, zero-sum game, or generic $2 \times m$ game. Let $\{a(t)\}_{t \geq 1}$ be a Computationally Efficient Sampled FP process, and assume A.2 holds. Then $\lim_{t \rightarrow \infty} d(q(t), NE) = 0$ (a.s.).*

Proof: We will prove the result by showing that there exists a sequence $\{\epsilon_t\}_{t \geq 1}$ with $\lim_{t \rightarrow \infty} \epsilon_t = 0$ such that $U_i(a_i(t+1), q_{-i}(t)) \geq \max_{\alpha_i \in A_i} U_i(\alpha_i, q_{-i}(t)) - \epsilon_t$. By [20], Corollary 5, this is sufficient to guarantee $\lim_{t \rightarrow \infty} d(q(t), NE) = 0$.

Since, by (6), $a_i(t+1) \in \arg \max_{\alpha_i \in A_i} \widehat{U}_i(\alpha_i, t)$, it is sufficient to show that, for every $i \in N$ and every $\alpha_i \in A_i$,

$$|\widehat{U}_i(\alpha_i, t) - U_i(\alpha_i, q_{-i}(t))| \rightarrow 0 \text{ as } t \rightarrow \infty \text{ a.s.} \quad (9)$$

(Note that the individual action spaces A_i are finite.) We will show this by invoking the result of Lemma 2 (see appendix). To that end fix $i \in N$ and $\alpha_i \in A_i$, and let $X(t) :=$

⁶Additional insight may be gained by considering the dynamical systems approach to stochastic approximations (e.g. [33]) which allows one to study the behavior of certain discrete-time processes by analyzing an associated differential equation. In such an analysis, an estimation rule such as (5) is often considered as a two time-scale system [26], [34], with the ODE associated with the estimation rule operating at a *faster* rate than ODE associated with the mixed utility process. In the asymptotic analysis of such systems, the slower process may often be treated as effectively static compared to the faster process. We note, however, that the proofs of our results rely primarily on self contained martingale-type arguments, rather than invoking results from dynamical systems based treatment of stochastic approximation literature.

$U_i(\alpha_i, a_{-i}^*(t))$, $t \geq 1$, $\mu(t) := U_i(\alpha_i, q_{-i}(t))$, $t \geq 1$, and $\hat{\mu}(t) := \hat{U}_i(\alpha_i, t)$, $t \geq 0$. For $t \geq 0$, let $\mathcal{F}_t := \sigma(\{q_{-i}(s)\}_{s=1}^{t+1})$. Note that $\mu(t)$ is \mathcal{F}_{t-1} -measurable and that $E(X(t)|\mathcal{F}_{t-1}) = \mu(t)$.

In order to invoke Lemma 2 it is sufficient to show that

$$\left(\frac{1}{\rho(t)} - 1\right) (\mu(t) - \mu(t-1)) \rightarrow 0. \quad (10)$$

Let $M := \max_{q'_{-i}, q''_{-i} \in \Delta_{-i}} \|q'_{-i} - q''_{-i}\|$, and note that by (7), $\|q_{-i}(t) - q_{-i}(t-1)\| \leq \frac{M}{t}$. The utility function U_i is multilinear, and hence Lipschitz continuous, so there exists a constant K such that

$$\begin{aligned} |\mu(t) - \mu(t-1)| &= |U_i(\alpha_i, q_{-i}(t)) - U_i(\alpha_i, q_{-i}(t-1))| \\ &\leq K \|q_{-i}(t) - q_{-i}(t-1)\| \leq KM/t. \end{aligned}$$

This, together with **A.2**, implies that (10) holds.

Thus, $X(t)$, $\mu(t)$, $\hat{\mu}(t)$, and \mathcal{F}_t as defined above fit the template of Lemma 2. By Lemma 2, $|\hat{U}_i(\alpha_i, t) - U_i(\alpha_i, q_{-i}(t))| = |\hat{\mu}(t) - \mu(t)| \rightarrow 0$ as $t \rightarrow \infty$, verifying that (9) holds. ■

VI. SIMULATION RESULTS

In order to demonstrate the computational properties of CESFP and Sampled FP in large games, we simulated both algorithms in a simple traffic routing scenario. Let $N = \{1, \dots, n\}$ denote a finite set of drivers (or players). Drivers share a common starting point and a common destination and may travel on one of 50 parallel routes. Let the set of routes be denoted by R , and let the action space of player i be given by $Y_i = R$, $\forall i$. Let $\sigma_r(y)$ denote the number of drivers on route r given the joint strategy y . Each route $r \in R$ has an associated cost function $c_r : \mathbb{N} \rightarrow \mathbb{R}$ signifying the delay experienced on route r given the number of drivers using the route. Let the utility function of player i be given by $u_i(y) := -c_{y_i}(\sigma_{y_i}(y))$. We note that this game is an instance of a congestion game—a known subset of potential games.

We simulated Sampled FP and CESFP in this routing scenario with 1000 drivers. In the simulation, Sampled FP used a sample rate of $k_t = \lfloor t^{.6} \rfloor$ samples per round and CESFP used the parameter $\rho(t) = t^{-.6}$, $\forall t$. Figure 1(a) shows the wall clock running time through iteration t for each algorithm.

Figure 1(b) shows a logarithmic plot of the expected total travel time if the current-iteration




Fig. 1. (a) Wall clock evaluation time for each algorithm; and (b) Expected total travel time for the mixed strategy $q(t)$.

empirical distribution $q(t)$ were to be used as the joint mixed strategy. While a Nash equilibrium of the traffic routing game does not necessarily minimize total travel time, the trend shown in Figure 1(b) is consistent with convergence of $q(t)$ to NE, and suggests a comparable convergence rate (per repeated-play iteration) for both algorithms.

VII. CONCLUSIONS

The classical Fictitious Play (FP) algorithm can be prohibitively difficult to implement in games with many players. Sampled FP [9]—a Monte-Carlo based variant of FP—has previously been proposed as a method of mitigating computational complexity in large-scale implementations of FP. Though Sampled FP does achieve mitigations in complexity, it suffers from the drawback that the number of samples that must be gathered in each stage of the algorithm grows without bound.

The paper proposed Computationally Efficient Sampled FP (CESFP)—a variant of Sampled FP that requires only one sample to be drawn per stage of the algorithm. CESFP is shown to achieve Nash equilibrium learning in the same sense as FP. A simulation example was used to demonstrate the computational properties of CESFP compared to Sampled FP. The simulation example used a game with fairly simple structural properties. An interesting future research direction may be to study the relative empirical performance of Sampled FP and CESFP in games with more complex structure (e.g. [10]).

APPENDIX

Lemma 1. Let $\{x_t\}_{t \geq 0}$ satisfy $x_t \rightarrow x$ as $t \rightarrow \infty$. Let $\{\rho_t\}_{t \geq 0}$ satisfy $0 < \rho_t \leq 1$ and $\sum_{t \geq 0} \rho_t = \infty$. Then, the sequence $\{y_t\}_{t \geq 0}$ given by $y_t = (1 - \rho_t)y_{t-1} + \rho_t x_t$, $t \geq 1$, satisfies $y_t \rightarrow x$ as $t \rightarrow \infty$.

Proof: The result follows from Toeplitz's lemma [36]. \blacksquare

Lemma 2. Let $\{\rho_t\}_{t \geq 1}$ satisfy $0 < \rho_t \leq 1$, $\sum_{t \geq 1} \rho_t = \infty$ and $\sum_{t \geq 1} \rho_t^2 < \infty$. Let $\{\mathcal{F}_t\}_{t \geq 1}$ be a filtration and let $\{X_t\}_{t \geq 1}$ be a sequence of bounded random variables, adapted to the filtration, say, $|X_t| \leq B$. Let $\mu_t = \mathbb{E}(X_t | \mathcal{F}_{t-1})$ and assume that $\left(\frac{1}{\rho_t} - 1\right)(\mu_t - \mu_{t-1}) \rightarrow 0$ almost surely. Then, the sequence of random variables $\{\hat{\mu}_t\}_{t \geq 0}$ given by $\hat{\mu}_t = (1 - \rho_t)\hat{\mu}_{t-1} + \rho_t X_t$, $t \geq 1$, satisfies $|\hat{\mu}_t - \mu_t| \rightarrow 0$ almost surely.

Proof: Subtracting μ_t from both sides of $\hat{\mu}_t = (1 - \rho_t)\hat{\mu}_{t-1} + \rho_t X_t$ gives $E_t = (1 - \rho_t)E_{t-1} + \rho_t \left(X_t - \mu_t + \left(\frac{1}{\rho_t} - 1\right)\delta_t\right)$, where $E_t := \hat{\mu}_t - \mu_t$ and $\delta_t := \mu_{t-1} - \mu_t$, $\delta_0 := 0$.

Introduce the \mathcal{F}_t -adapted sequences, for $t \geq 1$

$$\begin{aligned} F_t &= (1 - \rho_t)F_{t-1} + \rho_t(X_t - \mu_t), & F_1 &= E_1 \\ G_t &= (1 - \rho_t)G_{t-1} + \rho_t \left(\frac{1}{\rho_t} - 1\right)\delta_t, & G_1 &= 0, \end{aligned}$$

and note that $E_t = F_t + G_t$. We will now show that $F_t \rightarrow 0$ and $G_t \rightarrow 0$ almost surely.

By assumption, $\left(\frac{1}{\rho_t} - 1\right)\delta_t \rightarrow 0$ almost surely. Lemma 1 applied to $\{G_t\}_{t \geq 1}$ gives $G_t \rightarrow 0$ almost surely.

On the other hand,

$$\begin{aligned} \mathbb{E}(F_t^2 | \mathcal{F}_{t-1}) &= (1 - \rho_t)^2 F_{t-1}^2 + \rho_t^2 \mathbb{E}((X_t - \mu_t)^2 | \mathcal{F}_{t-1}) \\ &\leq (1 - \rho_t)^2 F_{t-1}^2 + \rho_t^2 4B^2 \\ &= (1 + \rho_t^2) F_{t-1}^2 - 2\rho_t F_{t-1}^2 + \rho_t^2 4B^2. \end{aligned}$$

Since $\sum_{t \geq 0} \rho_t^2 < \infty$, from the Robbins-Monro Lemma [37] we conclude that, almost surely, $\{F_t^2\}_{t \geq 1}$ converges and $\sum_{t \geq 0} \rho_t F_{t-1}^2 < \infty$: these two properties imply $F_t \rightarrow 0$ almost surely. \blacksquare

Lemma 3. Let $\{\rho_t\}_{t \geq 1}$ be such that $0 < \rho_t \leq 1$ and $\lim_{t \rightarrow \infty} \frac{1}{t\rho_t} = 0$. Then $\sum_{t \geq 1} \rho_t = \infty$.

Proof: We claim there exists $c, T > 0$ such that $\rho_t \geq c\frac{1}{t}$ for all $t \geq T$. If this were not so, then for every $c > 0$ there would hold $\rho_t < c\frac{1}{t}$ infinitely often, which would imply $\frac{1}{t\rho_t} > \frac{1}{c}$ infinitely often—contradicting the hypothesis that $\lim_{t \rightarrow \infty} \frac{1}{t\rho_t} = 0$.

Thus, there exists $c, T > 0$ such that $\rho_t \geq c\frac{1}{t} \forall t \geq T$, and hence $\sum_{t \geq 1} \rho_t \geq \sum_{t \geq T} \rho_t \geq c \sum_{t \geq T} \frac{1}{t} = \infty$. \blacksquare

REFERENCES

- [1] J. R. Marden and J. S. Shamma, "Game theory and distributed control," *Handbook of game theory*, vol. 4, 2012.
- [2] W. Saad, Z. Han, H. V. Poor, and T. Basar, "Game-theoretic methods for the smart grid: An overview of microgrid systems, demand-side management, and smart grid communications," *Signal Process. Mag., IEEE*, vol. 29, no. 5, pp. 86–105, 2012.
- [3] J. R. Marden, S. Ruben, and L. Pao, "A model-free approach to wind farm control using game theoretic methods," *IEEE Transactions on Control Systems Technology*, vol. 21, no. 4, pp. 1207–1214, 2013.

- [4] B. Wang, Y. Wu, and K. R. Liu, "Game theory for cognitive radio networks: An overview," *Computer networks*, vol. 54, no. 14, pp. 2537–2561, 2010.
- [5] G. Arslan, J. R. Marden, and J. S. Shamma, "Autonomous vehicle-target assignment: A game-theoretical formulation," *Journal of Dynamic Systems, Measurement, and Control*, vol. 129, no. 5, pp. 584–596, 2007.
- [6] B. Yang and M. Johansson, "Distributed optimization and games: A tutorial overview," in *Networked Control Systems*. Springer, 2010, pp. 109–148.
- [7] N. Li and J. R. Marden, "Designing games for distributed optimization," in *Proceedings of the 50th IEEE Conference on Decision and Control*, 2011.
- [8] J. R. Marden, H. P. Young, and L. Y. Pao, "Achieving Pareto optimality through distributed learning," in *IEEE 51st Annual Conference on Decision and Control*. IEEE, 2012, pp. 7419–7424.
- [9] T. J. Lambert, M. A. Epelman, and R. L. Smith, "A fictitious play approach to large-scale optimization," *Operations Research*, vol. 53, no. 3, pp. 477–489, 2005.
- [10] A. Garcia, D. Reaume, and R. L. Smith, "Fictitious play for finding system optimal routings in dynamic traffic networks," *Transportation Research Part B: Methodological*, vol. 34, no. 2, pp. 147–156, 2000.
- [11] G. W. Brown, "Iterative Solutions of Games by Fictitious Play" In *Activity Analysis of Production and Allocation*, T. Coopmans, Ed. New York: Wiley, 1951.
- [12] D. Fudenberg and D. K. Levine, *The theory of learning in games*. MIT press, 1998, vol. 2.
- [13] H. P. Young, *Strategic learning and its limits*. Oxford University Press, 2004.
- [14] D. Monderer and L. S. Shapley, "Fictitious play property for games with identical interests," *Journal of Economic Theory*, vol. 68, no. 1, pp. 258–265, 1996.
- [15] —, "Potential Games," *Games and Economic Behavior*, vol. 14, no. 1, pp. 124–143, 1996.
- [16] M. Epelman, A. Ghate, and R. L. Smith, "Sampled fictitious play for approximate dynamic programming," *Computers & Operations Research*, vol. 38, no. 12, pp. 1705–1718, 2011.
- [17] A. Ghate, S.-F. Cheng, S. Baumert, D. Reaume, D. Sharma, and R. L. Smith, "Sampled fictitious play for multi-action stochastic dynamic programs," *IIE Transactions*, vol. 46, no. 7, pp. 742–756, 2014.
- [18] T. J. Lambert and H. Wang, "Fictitious play approach to a mobile unit situation awareness problem," *Univ. Michigan, Tech. Rep*, 2003.
- [19] E. Sisikoglu, M. A. Epelman, and R. L. Smith, "A sampled fictitious play based learning algorithm for infinite horizon markov decision processes," in *Simulation Conference, Proceedings of the 2011 Winter*. IEEE, 2011, pp. 4086–4097.
- [20] D. S. Leslie and E. Collins, "Generalised weakened fictitious play," *Games and Economic Behavior*, vol. 56, no. 2, pp. 285–298, 2006.
- [21] K. Dabcevic, A. Betancourt, L. Marcenaro, and C. S. Regazzoni, "A fictitious play-based game-theoretical approach to alleviating jamming attacks for cognitive radios," in *Acoustics, Speech and Signal Processing (ICASSP), IEEE International Conference on*. IEEE, 2014, pp. 8158–8162.
- [22] —, "Intelligent cognitive radio jamming—a game-theoretical approach," *EURASIP Journal on Advances in Signal Processing*, vol. 2014, no. 1, 2014.
- [23] B. Swenson, S. Kar, and J. Xavier, "A computationally efficient implementation of fictitious play in a distributed setting," 2015, accepted for publication, EUSIPCO.
- [24] J. R. Marden, G. Arslan, and J. S. Shamma, "Joint strategy fictitious play with inertia for potential games," *IEEE Transactions on Automatic Control*, vol. 54, no. 2, pp. 208–220, 2009.

- [25] A. C. Chapman, D. S. Leslie, A. Rogers, and N. R. Jennings, “Convergent learning algorithms for unknown reward games,” *SIAM Journal on Control and Optimization*, vol. 51, no. 4, pp. 3154–3180, 2013.
- [26] D. S. Leslie, E. Collins *et al.*, “Convergent multiple-timescales reinforcement learning algorithms in normal form games,” *The Annals of Applied Probability*, vol. 13, no. 4, pp. 1231–1251, 2003.
- [27] G. Arslan and J. S. Shamma, “Distributed convergence to Nash equilibria with local utility measurements,” in *Proceedings of the 43rd IEEE Conference on Decision and Control*, vol. 2, 2004, pp. 1538 – 1543.
- [28] J. R. Marden, H. P. Young, G. Arslan, and J. S. Shamma, “Payoff-based dynamics for multi-player weakly acyclic games,” *SIAM Journal on Control and Optimization*, vol. 48, no. 1, pp. 373–396, 2009.
- [29] F. Germano and G. Lugosi, “Global Nash convergence of Foster and Young’s regret testing,” *Games and Economic Behavior*, vol. 60, no. 1, pp. 135–154, 2007.
- [30] B. S. Pradelski and H. P. Young, “Learning efficient Nash equilibria in distributed systems,” *Games and Economic Behavior*, vol. 75, no. 2, pp. 882–879, 2012.
- [31] B. Swenson, S. Kar, and J. Xavier, “Empirical centroid fictitious play: An approach for distributed learning in multi-agent games,” *IEEE Transactions on Signal Processing*, 2015, In Press. Initial Submission: Apr. 2013. Accepted: Mar. 2015. [Online]: <http://arxiv.org/pdf/1304.4577v2.pdf>.
- [32] D. Williams, *Probability with martingales*. Cambridge university press, 1991.
- [33] V. S. Borkar, *Stochastic approximation*. Cambridge University Press, 2008.
- [34] —, “Stochastic approximation with two time scales,” *Systems & Control Letters*, vol. 29, no. 5, pp. 291–294, 1997.
- [35] U. Berger, “Fictitious play in $2 \times n$ games,” *Journal of Economic Theory*, vol. 120, no. 2, pp. 139–154, 2005.
- [36] M. Loève, *Probability Theory I*. Springer-Verlag, 1977.
- [37] H. Robbins and D. Siegmund, “A convergence theorem for non negative almost supermartingales and some applications,” in *Herbert Robbins Selected Papers*. Springer, 1985, pp. 111–135.